



PDSF

NERSC's Production Linux Cluster

Craig E. Tull
HCG/NERSC/LBNL
MRC Workshop
LBNL - *March 26, 2002*



Outline



- **People: Shane Canon (Lead), Cary Whitney, Tom Langley, Iwona Sakredja (Support)**
(Tom Davis, Tina Declerk, John Milford, others)
- **Present**
 - What is PDSF?
 - Scale, HW & SW Architecture, Business & Service Models, Science Projects
- **Past**
 - Where did PDSF come from?
 - Origins, Design, Funding Agreement
- **Future**
 - How does PDSF relate to the MRC initiative?



PDSF - Production Cluster



- **PDSF - Parallel Distributed Systems Facility**
—HENP community
 - Specialized needs/Specialized requirements
- **Our mission is to provide the most effective distributed computer cluster possible that is suitable for experimental HENP applications.**
- **Architecture tuned for “embarrassingly parallel” applications**
- **AFS access, and access to HPSS for mass storage**
- **High speed (Gigabit Ethernet) access to HPSS system and to Internet Gateway**
- **<http://pdsf.nersc.gov/>**



PDSF Photo



LAWRENCE BERKELEY NATIONAL LABORATORY

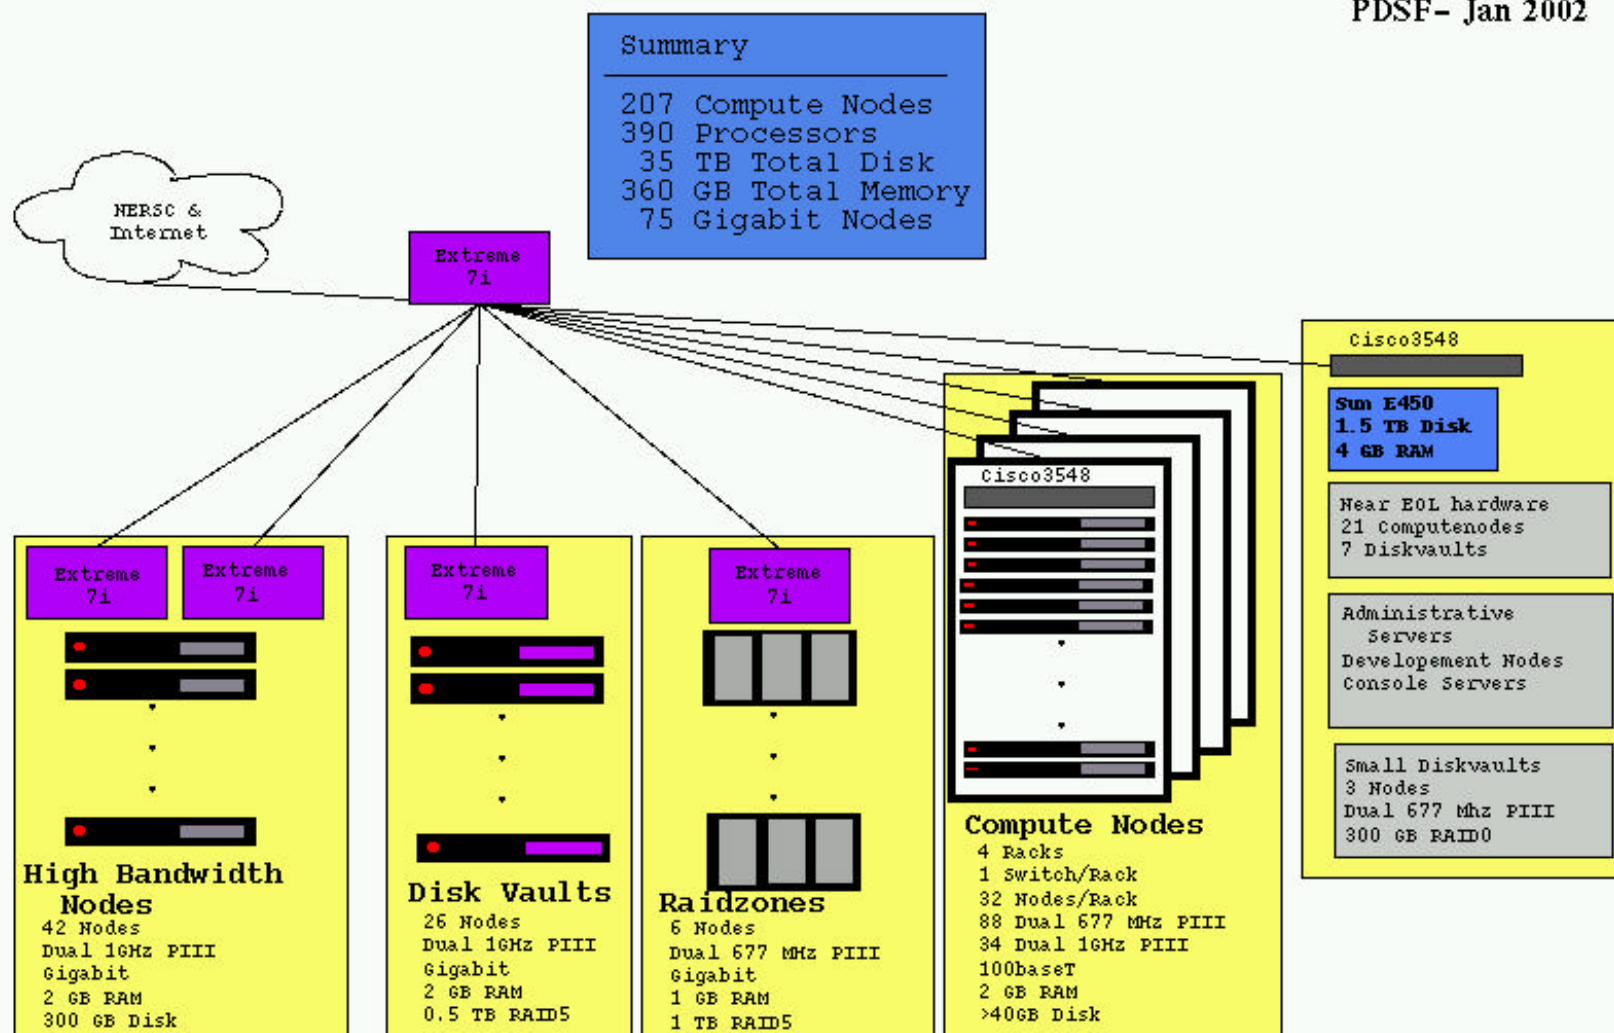
CETull@lbl.gov - PDSF - Nersc's Production Linux Cluster (26mar02 - MRC WkShp @ LBNL)



PDSF Overview



PDSF- Jan 2002





Component HW Architecture



- By minimizing diversity within the system, we maximize uptime and minimize sys-admin overhead.
- Buying Computers by the Slice (Plant & Prune)
 - Buy large, homogeneous batches of HW.
 - Each large purchase of HW can be managed as a "single unit" composed of interchangeable parts.
 - Each slice has limited lifespan - **NO MAINT. \$**
- Critical vs. non-Critical Resources
 - Non-critical compute nodes can fail without stopping analysis.
 - Critical data vaults can be trivially interchanged (with transfer of disks) with compute nodes.
- Uniform environment means that software & security problems can be solved by "Reformat & Reload".



HW Arch.: 4 types of Nodes



- **Interactive**
 - 8 Intel Linux (RedHat 6.2)
 - More memory, Fast, interactive logon, serve batch jobs when idle
- **Batch (Normal & High Bandwidth)**
 - ~400 Intel Linux CPUs (RedHat 6.2)
 - LSF: Short, Medium, Long, & Custom Queues
- **Data Vault**
 - Large Shared (NFS) Disk Arrays - 25 TB
 - High Network (GigE) Connectivity to compute nodes & HPSS, Data-Transfer Jobs Only
- **Administrative**
 - Home diskspace server, AFS servers, License servers, Database servers, time servers, etc.



PDSF Global SW



- All SW necessary to HENP data analysis & simulation is available and maintained at current revs
- Solaris Software
 - AFS, CERN libs, CVS, Modules, Objectivity, Omnibroker, Orbix, HPSS pFTP, ssh, LSF, PVM, Framereader, Sun Workshop Suite, Sun's Java Dev. Toolkit, Veritas, Python, etc.
- Linux Software
 - AFS, CERN libs, CVS, Modules, Omnibroker, ssh, egcs, KAI C++, Portland Group F77/F90/HPF/C/C++, LSF, HPSS pFTP, Objectivity, ROOT, Python, etc.
- Specialized Software (Experiment Maintained)
 - ATLAS, CDF, D0, DPSS, E895, STAR, etc.





Experiment/Project SW



- **Principal: Allow diverse groups/projects to operate on all nodes without interfering with others.**
- **Modules:**
 - **Allows individuals/groups to chose appropriate SW versions at login (Version migration dictated by experiment, not system.).**
- **Site independence:**
 - **PDSF personnel have been very active in helping "portify" code (STAR, ATLAS, CDF, ALICE).**
 - **Direct benefit to project Regional Centers & institutions.**
- **Specific kernel/libc dependency of project SW is only case where interference is an issue (None now.).**
- **LSF extensible to allow incompatible differences.**



Batch Queuing System



- **LSF 4.x - Load Sharing Facility**
 - Solaris & Linux
 - Tremendous leverage from NERSC (158 "free", aggressive license negotiations & price savings)
 - Very good user & admin experiences
 - Fair share policy in use
 - Can easily sustain >95% load.



Administrative SW/Tools



- **Monitoring tools - Batch usage, node & disk health, etc.**
 - developed at PDSF to insure smooth operation & assure contributing clients
- **HW tracking & location (mySQL + ZOPE)**
 - 800 drives, ~300 boxes, HW failures/repairs, etc.
 - developed at PDSF out of absolute necessity
- **System & Package consistency**
 - developed at PDSF
- **System installation**
 - kickstart - ~3 min.s/node
- **System Security - No Known Security Breaches**
 - TCP Wrappers & ipchains, NERSC Monitoring, no clear-text passwords, security patches, crack, etc.



PDSF Business Model



- Users contribute directly to cluster through hardware purchases. The size of the contribution determines the fraction of resources that are **guaranteed**.
- NERSC provides facilities and administrative resources (up to pre-agreed limit).
- User share of resource guaranteed at 100% for 2 years (warranty), then depreciates 25% per year (hardware lifespan).
- PDSF scale has reached the point where some FTE resources must be funded.
 - #Admins \propto Size of System (eg. box count)
 - #Support \propto Size of User Community (eg. #groups & #users & diversity)



PDSF Service Model



- **Less than 24/7, but more than Best Effort.**
- **Support (1 FTE):**
 - USG supported web-based trouble tickets.
 - Response during business hours.
 - Performance matrix.
 - Huge load right now Û Active, Large community
- **Admin (3 FTE):**
 - NERSC operations monitoring (24/7)
 - Critical vs. non-critical resources
 - Non-critical: (eg. batch nodes) Best effort
 - Critical: (eg. servers, DVs) Fastest response



History of PDSF Hardware




- Arrived from SSC (RIP), May 1997
- October '97 - "Free" HW from SSC
 - 32 SUN Sparc 10, 32 HP 735
 - 2 SGI data vaults
- 1998 - NERSC seeding & initial NSD updates
 - **added** 12 Intel (E895), SUN E450, 8 dual-cpu Intel (NSD/STAR), 16 Intel (NERSC), 500 GB network disk (NERSC)
 - **subtracted** SUN, HP, 160 GB SGI data vaults
- 1999 - Present - Full Plant & Prune
 - HENP Contributions: STAR, E871, SNO, ATLAS, CDF, E891, others
- March 2002
 - 240 Intel Compute Nodes (390 CPUs)
 - 8 Intel Interactive Nodes (dual 996MHz PIII, 2GB RAM, 55GB scratch)
 - 49 Data Vaults: 25 TB of shared disk
 - Totals: 570K MIPS, 35 TB disk



PDSF Growth



- Future
 - Continued support of STAR & HENP
 - Continue to seek out new groups
 - Look for interest outside of HENP community
 - Primarily Serial workloads
 - Stress benefits of using a shared resource
 - Let scientist focus on science and system administrator run computer systems





PDSF Major/Active Users






- **Collider Facilities**
 - STAR/RHIC - Largest user
 - CDF/FNAL
 - ATLAS/CERN
 - E871/FNAL
- **Neutrino Experiments**
 - Sudbury Neutrino Observatory (SNO)
 - KamLAND
- **Astrophysics**
 - Deep Search
 - Super Nova Factory
 - Large Scale Structure
 - AMANDA



Solenoidal Tracking At RHIC (STAR)



- Experiment at the RHIC accelerator in BNL
- Over 400 scientists and engineers from 33 institutions in 7 countries
- PDSF primarily intended to handle simulation workload
- PDSF has increasingly been used for general analysis




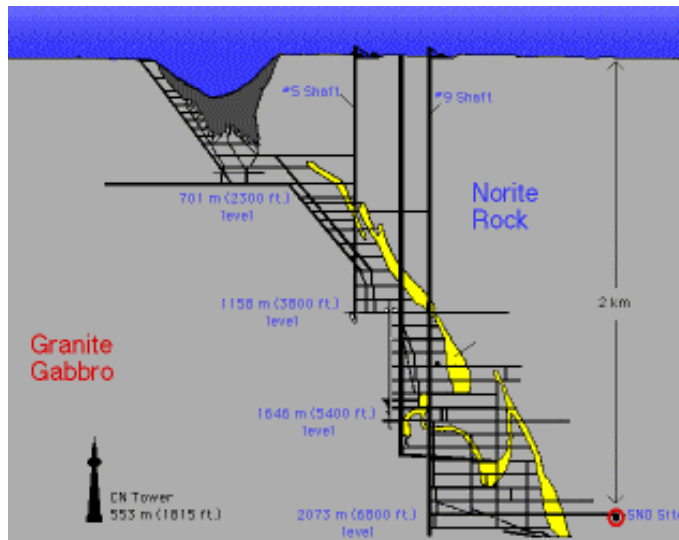
LAWRENCE BERKELEY NATIONAL LABORATORY



Sudbury Neutrino Observatory (SNO)



- Located in a mine in Ontario Canada
- Heavy water neutrino detector
- SNO has over 100 collaborators at 11 institutions





Results from SNO



- Recently confirmed results from Super-K that neutrinos have mass.
- PDSF specifically mentioned in results

Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^8B solar neutrinos at the Sudbury Neutrino Observatory

Q.R. Ahmad¹⁵, R.C. Allen¹¹, T.C. Andersen¹², J.D. Anglin⁷, G. Bühler¹¹, J.C. Barton^{13†}, E.W. Beier¹⁴,
M. Bercovitch⁷, J. Bigu⁴, S. Biller¹³, R.A. Black¹³, I. Blevis³, R.J. Boardman¹³, J. Boger², E. Bonvin⁹,
M.G. Boulay⁹, M.G. Bowler¹³, T.J. Bowles⁶, S.J. Brice^{6,13}, M.C. Browne¹⁵, T.V. Bullard¹⁵, T.H. Burritt^{15,6},
K. Cameron¹², J. Cameron¹³, Y.D. Chan⁵, M. Chen⁹, H.H. Chen^{11*}, X. Chen^{5,13}, M.C. Chon¹², B.T. Cleveland¹³,
E.T.H. Clifford^{9,1}, J.H.M. Cowan⁴, D.F. Cowen¹⁴, G.A. Cox¹⁵, Y. Dai⁹, X. Dai¹³, F. Dalnoki-Veress³,
W.F. Davidson⁷, P.J. Doe^{15,11,6}, G. Doucas¹³, M.R. Dragowsky^{6,5}, C.A. Duba¹⁵, F.A. Duncan⁹, J. Dunmore¹³,
E.D. Earle^{9,1}, S.R. Elliott^{15,6}, H.C. Evans⁹, G.T. Ewan⁹, J. Farine³, H. Fergani¹³, A.P. Ferraris¹³, R.J. Ford⁹,
M.M. Fowler⁶, K. Frame¹³, E.D. Frank¹⁴, W. Frati¹⁴, J.V. Germani^{15,6}, S. Gil¹⁰, A. Goldschmidt⁶, D.R. Grant³,
R.L. Hahn², A.L. Hallin⁹, E.D. Hallman⁴, A. Hamer^{6,9}, A.A. Hamian¹⁵, R.U. Haq⁴, C.K. Hargrove³, P.J. Harvey⁹,
R. Hazama¹⁵, R. Heaton⁹, K.M. Heeger¹⁵, W.J. Heintzelman¹⁴, J. Heise¹⁰, R.L. Helmer^{10†}, J.D. Hepburn^{9,1},
H. Heron¹³, J. Hewett⁴, A. Hime⁶, M. Howe¹⁵, J.G. Hykawy⁴, M.C.P. Isaac⁵, P. Jagam¹², N.A. Jelley¹³, C. Jillings⁹,
G. Jonkmans^{4,1}, J. Karn¹², P.T. Keener¹⁴, K. Kirch⁶, J.R. Klein¹⁴, A.B. Knox¹³, R.J. Komar^{10,9}, R. Kouzes⁸,
T. Kutter¹⁰, C.C.M. Kyba¹⁴, J. Law¹², I.T. Lawson¹², M. Lay¹³, H.W. Lee⁹, K.T. Lesko⁵, J.R. Leslie⁹, I. Levine³,
W. Locke¹³, M.M. Lowry⁸, S. Luoma⁴, J. Lyon¹³, S. Majerus¹³, H.B. Mak⁹, A.D. Marino⁵, N. McCauley¹³,
A.B. McDonald^{9,8}, D.S. McDonald¹⁴, K. McFarlane³, G. McGregor¹³, W. McLatchie⁹, R. Meijer Drees¹⁵, H. Mes³,
C. Mifflin³, G.G. Miller⁶, G. Milton¹, B.A. Moffat⁹, M. Moorhead^{13,5}, C.W. Nally¹⁰, M.S. Neubauer¹⁴,
E.M. ...

LAWRENCE BERKELEY NATIONAL LABORATORY



PDSF Job Matching



- Linux clusters are already playing a large role in HENP and Physics simulations and analysis.
- Beowulf systems may be in-expensive, but can require lots of time to administer
- Serial Jobs:
 - Perfect match (up to ~2GB RAM+2GB SWAP)
- MPI Jobs:
 - Some projects using MPI (Large Scale Structures, Deep Search).
 - Not** low latency, small messages
- "Real" MPP (eg. Myricom)
 - Not currently possible. Could be done, but entails significant investment of time & money.
 - LSF handles MPP jobs (already configured).



How does PDSF relate to MRC?



- **Emulation or Expansion ("Model vs. Real Machine")**
 - If job characteristics & resources match.
- **Emulation:**
 - Adopt appropriate elements of HW & SW Arch.s, Business & Service Models.
 - Steal appropriate tools.
 - Customize to eg. non-PDSF like job load.
- **Expansion:**
 - Directly contribute to PDSF resource.
 - Pros:
 - Faster ramp-up, stable environment, try before buy, larger pool of resources (better utilization), leverage existing expertise/infrastructure
 - Cons:
 - Not currently MPP tuned, ownership issues.